Integer Programming and Combinatorial Optimization (Bixby et al. eds.), Springer LNCS 1412, 1998, 195-201

The Maximum Traveling Salesman Problem under Polyhedral Norms

Alexander Barvinok[1], David S. Johnson[2], Gerhard J. Woeginger[3], and Russell Woodroofe[4]

[1] University of Michigan, Dept. of Mathematics,
Ann Arbor, MI 48109-1009, USA
barvinok@math.lsa.umich.edu
Supported by an Alfred P. Sloan Research Fellowship and NSF grant DMS 9501129
[2] AT&T Labs – Research, Room C239,
Florham Park, NJ 07932-0971, USA
dsj@research.att.com
[3] Institut für Mathematik, TU Graz,
Steyrergasse 30, A-8010 Graz, Austria
gwoegi@opt.math.tu-graz.ac.at
Supported by the START program Y43-MAT of the Austrian Ministry of Science
[4] University of Michigan, Dept. of Mathematics,
Ann Arbor, MI 48109-1009, USA
Supported by the NSF through the REU Program

Abstract. We consider the traveling salesman problem when the cities are points in $\boldsymbol{R}^d$ for some fixed d and distances are computed according to a polyhedral norm. We show that for any such norm, the problem of finding a tour of *maximum* length can be solved in polynomial time. If arithmetic operations are assumed to take unit time, our algorithms run in time $O(n^{f-2} \log n)$, where f is the number of facets of the polyhedron determining the polyhedral norm. Thus for example we have $O(n^2 \log n)$ algorithms for the cases of points in the plane under the Rectilinear and Sup norms. This is in contrast to the fact that finding a *minimum* length tour in each case is NP-hard.

1 Introduction

In the *Traveling Salesman Problem* (TSP), the input consists of a set C of *cities* together with the distances $d(c, c')$ between every pair of distinct cities $c, c' \in C$. The goal is to find an ordering or *tour* of the cities that minimizes (Minimum TSP) or maximizes (Maximum TSP) the total tour length. Here the length of a tour $c_{\pi(1)}, c_{\pi(2)}, \ldots, c_{\pi(n)}$ is

$$\sum_{i=1}^{n-1} d(c_{\pi(i)}, c_{\pi(i+1)}) + d(c_{\pi(n)}, c_{\pi(1)}).$$

Of particular interest are *geometric* instances of the TSP, in which cities correspond to points in $\boldsymbol{R}^d$ for some $d \geq 1$, and distances are computed according to some geometric norm. Perhaps the most popular norms are the Rectilinear, Euclidean, and Sup norms. These are examples of what is known as an "L^p norm" for $p = 1$, 2, and ∞. In general, the distance between two points $\mathbf{x} = (x_1, x_2, \ldots, x_d)$ and $\boldsymbol{y} = (y_1, y_2, \ldots, y_d)$ under the L^p norm, $p \geq 1$, is

$$d(\mathbf{x}, \boldsymbol{y}) = \left(\sum_{i=1}^{d} |x_i - y_i|^p \right)^{1/p}$$

with the natural asymptotic interpretation that distance under the L^∞ norm is

$$d(\mathbf{x}, \boldsymbol{y}) = \max_{i=1}^{d} |x_i - y_i| \,.$$

This paper concentrates on a second class of norms which also includes the Rectilinear and Sup norms, but can only approximate the Euclidean and other L^p norms. This is the class of *polyhedral norms*. Each polyhedral norm is determined by a *unit ball* which is a centrally-symmetric polyhedron $\boldsymbol{P}$ with the origin at its center. To determine $d(\mathbf{x}, \boldsymbol{y})$ under such a norm, first translate the space so that one of the points, say $\mathbf{x}$, is at the origin. Then determine the unique factor α by which one must rescale $\boldsymbol{P}$ (expanding if $\alpha > 1$, shrinking if $\alpha < 1$) so that the other point ($\boldsymbol{y}$) is on the boundary of the polyhedron. We then have $d(\mathbf{x}, \boldsymbol{y}) = \alpha$.

Alternatively, and more usefully for our purposes, we can view a polyhedral norm as follows. If $\boldsymbol{P}$ is a polyhedron as described above and has F facets, then F is divisible by 2 and there is a set $H_{\boldsymbol{P}} = \{\boldsymbol{h}_1, \ldots, \boldsymbol{h}_{F/2}\}$ of points in $\boldsymbol{R}^d$ such that $\boldsymbol{P}$ is the intersection of a collection of half-spaces determined by $H_{\boldsymbol{P}}$:

$$\boldsymbol{P} = \left(\bigcap_{i=1}^{F/2} \{\mathbf{x} : \ \mathbf{x} \cdot \boldsymbol{h}_i \leq 1\} \right) \cap \left(\bigcap_{i=1}^{F/2} \{\mathbf{x} : \ \mathbf{x} \cdot \boldsymbol{h}_i \geq -1\} \right)$$

Then we have

$$d(\mathbf{x}, \boldsymbol{y}) = \max \left\{ \Big| (\mathbf{x} - \boldsymbol{y}) \cdot \boldsymbol{h}_i \Big| : \ 1 \leq i \leq F/2 \right\}$$

Note that for the Rectilinear norm in the plane we can take $H_{\boldsymbol{P}} = \{(1,1), (-1,1)\}$ and for the Sup norm in the plane we can take $H_{\boldsymbol{P}} = \{(1,0), (0,1)\}$.

For the Minimum TSP on geometric instances, all of the key complexity questions have been answered. As follows from results of Itai, Papadimitriou, and Swarcfiter [5], the Minimum TSP is NP-hard for any fixed dimension d and any L^p or polyhedral norm. On the other hand, recent results of Arora [1,2] and Mitchell [8] imply that in all these cases a polynomial-time approximation scheme (PTAS) exists, i.e., a sequence of polynomial-time algorithms A_k, $1 \leq k < \infty$, where A_k is guaranteed to find a tour whose length is within a ratio of $1 + (1/k)$ of optimal.

The situation for geometric versions of the Maximum TSP is less completely resolved. Barvinok [3] has shown that once again polynomial-time approximation schemes exist for all fixed dimensions d and all L^p or polyhedral norms (and in a sense for *any* fixed norm; see [3]). Until now, however, the complexity of the optimization problems themselves when d is fixed has remained open: For no fixed dimension d and L^p or polyhedral norm was the problem of determining the maximum tour length known either to be NP-hard or to be polynomial-time solvable. In this paper, we resolve the question for polyhedral norms, showing that, in contrast to the case for the Minimum TSP, the Maximum TSP is solvable in polynomial time for any fixed dimension d and any polyhedral norm:

Theorem 1. *Let dimension d be fixed, and let $||\cdot||$ be a fixed polyhedral norm in $\mathbf{R}^d$ whose unit ball is a polyhedron $\boldsymbol{P}$ determined by a set of f facets. Then for any set of n points in $\mathbf{R}^d$, one can construct a traveling salesman tour of maximum length with respect to $||\cdot||$ in time $O(n^{f-2}\log n)$, assuming arithmetic operations take unit time.*

As an immediate consequence of Theorem 1, we get relatively efficient algorithms for the Maximum TSP in the plane under Rectilinear and Sup norms:

Corollary 1. *The Maximum TSP for points in $\mathbf{R}^2$ under the L^1 and L^∞ norms can be solved in $O(n^2\log n)$ time, assuming arithmetic operations take unit time.*

The restriction to unit cost arithmetic operations in Theorem 1 and Corollary 1 is made primarily to simplify the statements of the conclusions, although it does reflect that fact that our results hold for the *real number RAM* computational model. Suppose on the other hand that one assumes, as one typically must for complexity theory results, that the components of the vectors in $H_{\boldsymbol{P}}$ and the coordinates of the cities are all rationals. Let U denote the maximum absolute value of any of the corresponding numerators and denominators. Then the conclusions of the Theorem and Corollary hold with running times multipled by $n\log(U)$. If the components/coordinates are all integers with maximum absolute value U, the running times need only be multiplied by $\log(nU)$. For simplicity in the remainder of this paper, we shall stick to the model in which numbers can be arbitrary reals and arithmetic operations take unit time. The reader should have no trouble deriving the above variants.

The paper is organized as follows. Section 2 introduces a new special case of the TSP, the *Tunneling TSP*, and shows how the Maximum TSP under a polyhedral norm can be reduced to the Tunneling TSP with the same number of cities and $f/2$ tunnels. Section 3 sketches how the latter problem can be solved in $O(n^{f+1})$ time, a slightly weaker result than that claimed in Theorem 1. The details of how to improve this running time to $O(n^{f-2}\log n)$ will be presented in the full version of this paper, a draft of which is available from the authors. Section 4 concludes by describing some related results and open problems.

2 The Tunneling TSP

The *Tunneling TSP* is a special case of the Maximum TSP in which distances are determined by what we shall call a *tunnel system* distance function. In such a distance function we are given a set $T = \{t_1, t_2, \ldots, t_k\}$ of auxiliary objects that we shall call *tunnels*. Each tunnel is viewed as a bidirectional passage having a front and a back end. For each pair c, t of a city and a tunnel we are given real-valued *access distances* $F(c,t)$ and $B(c,t)$ from the city to the front and back ends of the tunnel respectively. Each potential tour edge $\{c, c'\}$ must pass through some tunnel t, either by entering the front end and leaving the back (for a distance of $F(c,t)+B(c',t)$), or by entering the back end and leaving the front (for a distance of $B(c,t)+F(c',t)$). Since we are looking for a tour of maximum length, we can thus define the distance between cities c and c' to be

$$d(c,c') = \max\Big\{F(c,t_i)+B(c',t_i),\ B(c,t_i)+F(c',t_i) : 1 \le i \le k\Big\}$$

Note that this distance function, like our geometric norms, is symmetric.

It is easy to see that Maximum TSP remains NP-hard when distances are determined by arbitrary tunnel system distance functions. However, for the case where $k = |T|$ is fixed and not part of the input, we will show in the next section that Maximum TSP can be solved in $O(n^{2k-1})$ time. We are interested in this special case because of the following lemma.

Lemma 1. *If $\|\cdot\|$ is a polyhedral norm determined by a set $H_{\boldsymbol{P}}$ of k vectors in $\boldsymbol{R}^d$, then for any set C of points in $\boldsymbol{R}^d$ one can in time $O(dk|C|)$ construct a tunnel system distance function with k tunnels that yields $d(\mathbf{c},\mathbf{c}') = \|\mathbf{c}-\mathbf{c}'\|$ for all $\mathbf{c},\mathbf{c}' \in C$.*

Proof. The polyhedral distance between two cities $\boldsymbol{c}, \boldsymbol{c}' \in \boldsymbol{R}^d$ is

$$\begin{aligned}\|\boldsymbol{c}-\boldsymbol{c}'\| &= \max\Big\{\Big|(\boldsymbol{c}-\boldsymbol{c}')\cdot\boldsymbol{h}_i\Big| :\ 1 \le i \le k\Big\}\\ &= \max\Big\{(\boldsymbol{c}-\boldsymbol{c}')\cdot\boldsymbol{h}_i,\ (\boldsymbol{c}'-\boldsymbol{c})\cdot\boldsymbol{h}_i :\ 1 \le i \le k\Big\}\end{aligned}$$

Thus we can view the distance function determined by $\|\cdot\|$ as a tunnel system distance function with set of tunnels $T = H_{\boldsymbol{P}}$ and $F(\boldsymbol{c},\boldsymbol{h}) = \boldsymbol{c}\cdot\boldsymbol{h}$, $B(\boldsymbol{c},\boldsymbol{h}) = -\boldsymbol{c}\cdot\boldsymbol{h}$ for all cities $\boldsymbol{c}$ and tunnels $\boldsymbol{h}$. □

3 An Algorithm for Bounded Tunnel Systems

This section is devoted to the proof of the following lemma, which together with Lemma 1 implies that the Maximum TSP problem for a fixed polyhedral norm with f facets can be solved in $O(n^{f+1})$ time.

Lemma 2. *If the number of tunnels is fixed at k, the Tunneling TSP can be solved in time $O(n^{2k+1})$, assuming arithmetic operations take unit time.*

Proof. Suppose we are given an instance of the Tunneling TSP with sets $C = \{c_1, \ldots, c_n\}$ and $T = \{t_1, \ldots, t_k\}$ of cities and tunnels, and access distances $F(c,t)$, $B(c,t)$ for all $c \in C$ and $t \in T$. We begin by transforming the problem to one about subset construction.

Let $G = (C \cup T, E)$ be an edge-weighted, bipartite multigraph with four edges between each city c and tunnel t, denoted by $e_i[c,t,X]$, $i \in \{1,2\}$ and $X \in \{B,F\}$. The weights of these edges are $w(e_i[c,t,F]) = F(c,t)$ and $w(e_i[c,t,B]) = B(c,t)$, $i \in \{1,2\}$. For notational convenience, let us partition the edges in E into sets $E[t,F] = \{e_i[c,t,F] : c \in C, i \in \{1,2\}\}$ and $E[t,B] = \{e_i[c,t,B] : c \in C, i \in \{1,2\}\}$, $t \in T$. Each tour for the TSP instance then corresponds to a subset E' of E that has $\sum_{e \in E'} w(e)$ equal to the tour length and satisfies

(T1) Every city is incident to exactly two edges in E'.
(T2) For each tunnel $t \in T$, $|E' \cap E[t,F]| = |E' \cap E[t,B]|$.
(T3) The set E' is connected.

To construct the multiset E', we simply represent each tour edge $\{c,c'\}$ by a pair of edges from E that connect in the appropriate way to the tunnel that determines $d(c,c')$. For example, if $d(c,c') = F(c,t) + B(c',t)$, and c appears immediately before c' when the tour is traversed starting from $c_{\pi(1)}$, then the edge (c,c') can be represented by the two edges $e_2[c,t,F]$ and $e_1[c',t,B]$. Note that there are enough (city,tunnel) edges of each type so that all tour edges can be represented, even if a given city uses the same tunnel endpoint for both its tour edges. Also note that if $d(c,c')$ can be realized in more than one way, the multiset E' will not be unique. However, any multiset E' constructed in this fashion will still have $\sum_{e \in E'} w(e)$ equal to the tour length.

On the other hand, any set E' satisfying (T1) – (T3) corresponds to one (or more) tours having length at least $\sum_{e \in E'} w(e)$: Let $T' \in T$ be the set of tunnels t with $|E' \cap E[t,F]| > 0$. Then $G' = (C \cup T', E')$ is a connected graph all of whose vertex degrees are even by (T1) – (T3). By an easy result from graph theory, this means that G' contains an Euler tour that by (T1) passes through each city exactly once, thus inducing a TSP tour for C. Moreover, by (T2) one can construct such an Euler tour with the additional property that if $e_i[c,t,x]$ and $e_j[c',t,y]$ are consecutive edges in this tour, then $x \neq y$, i.e, either $x = F, y = B$ or $x = B, y = F$. Thus we will have $w(e_i[c,t,x]) + w(e_j[c',t,y]) \leq d(c,c')$, and hence the length of the TSP tour will be at least $\sum_{e \in E'} w(e)$, as claimed.

Thus our problem is reduced to finding a maximum weight set of edges $E' \subseteq E$ satisfying (T1) – (T3). We will now sketch our approach to solving this latter problem; full details are available from the authors and will appear in the journal version of this paper. The basic idea is to divide into $O(n^{2k-2})$ subproblems, each of which can be solved in linear time. Each subcase corresponds to a choice of a degree sequence $\boldsymbol{d}$ for the tunnels, for which there are $O(n^{k-1})$ possibilities, and a choice, for that degree sequence, of a canonical-form sequence $\boldsymbol{s}$ of edges that connects together those tunnels that have positive degree, for which there are again $O(n^{k-1})$ possibilities.

Having chosen $\boldsymbol{d}$ and $\boldsymbol{s}$, what we are left with is a maximum weight bipartite b-matching problem: starting with a set consisting of the edges specified by $\boldsymbol{s}$, each tunnel end must have its degree augmented up to that specified by $\boldsymbol{d}$ and each city must have its degree augmented up to 2. This b-matching problem can be solved in $O(n^3)$ time by the standard technique that converts it to an assignment problem on an expanded graph. The overall running time for the algorithm is thus $O(n^{k-1}n^{k-1}n^3) = O(n^{2k+1})$, as claimed. □

In the full paper we show how two additional ideas enable us to reduce our running times to $O(n^{2k-2}\log n)$, as needed for the proof of Theorem 1. The first idea is to view each b-matching problem as a transportation problem with a bounded number of customer locations. This latter problem can be solved in linear time by combining ideas from [7, 4, 11]. The second idea is to exploit the similarities between the transportation instances we need to solve. Here a standard concavity result implies that one dimension of our search over degree sequences can be handled by a binary search. In the full paper we also discuss how the constants involved in our algorithms grow with k.

4 Conclusion

We have derived a polynomial time algorithm for the Maximum TSP when the cities are points in $\boldsymbol{R}^d$ for some fixed d and when the distances are measured according to some polyhedral norm. The complexity of the Maximum TSP with *Euclidean distances* and fixed d remains unsettled, however, even for $d = 2$. Although the Euclidean norm can be approximated arbitrarily closely by polyhedral norms (and hence Barvinok's result [3] that Maximum TSP has a PTAS), it is not itself a polyhedral norm.

A further difficulty with the Euclidean norm (one shared by both the Minimum and Maximum TSP) is that we still do not know whether the TSP is in NP under this norm. Even if all city coordinates are rationals, we do not know how to compare a tour length to a given rational target in less than exponential time. Such a comparison would appear to require us to evaluate a sum of n square roots to some precision, and currently the best upper bound known on the number of bits of precision needed to insure a correct answer remains exponential in n. Thus even if we were to produce an algorithm for the Euclidean Maximum TSP that ran in polynomial time when arithmetic operations (and comparisons) take unit time, it might not run in polynomial time on a standard Turing machine.

Another set of questions concerns the complexity of the Maximum TSP when d is *not* fixed. It is relatively easy to show that the problem is NP-hard for all L^p norms (the most natural norms that are defined for all $d > 0$). For the case of L^∞ one can use a transformation from Hamiltonian Circuit in which each edge is represented by a separate dimension. For the L^p norms, $1 \le p < \infty$, one can use a transformation from the Hamiltonian Circuit problem for cubic graphs, with a dimension for each *non*-edge. However, this still leaves open the question of whether there might exist a PTAS for any such norm when d is not fixed. Trevisan [10] has shown that the Minimum TSP is Max SNP-hard for any such norm, and so cannot have such PTAS's unless P = NP. We can obtain a similar result for the Maximum TSP under L^∞ by modifying our NP-hardness transformation so that the source problem is the Minimum TSP with all edge lengths in $\{1, 2\}$, a special case that was proved Max SNP-hard by Papadimitriou and Yannakakis [9]. The question remains open for L^p, $1 \leq p < \infty$, although we conjecture that these cases are Max SNP-hard as well.

Finally, we note that our results can be extended in several ways. For instance, one can get polynomial-time algorithms for asymmetric versions of the Maximum TSP in which distances are computed based on non-symmetric unit balls. Also, algorithmic approaches analogous to ours can be applied to geometric versions of other NP-hard maximization problems: For example, consider the *Weighted 3-Dimensional Matching Problem* that consists in partitioning a set of $3n$ elements into n triples of maximum total weight. The special case where the elements are points in $\boldsymbol{R}^d$ and where the weight of a triple equals the perimeter of the corresponding triangle measured according to some fixed polyhedral norm can be solved in polynomial time.

Acknowledgement. We thank Arie Tamir for helpful comments on a preliminary version of this paper, and in particular for pointing out that a speedup of $O(n^2)$ could be obtained by using the transportation problem results of [7] and [11]. Thanks also to Mauricio Resende and Peter Shor for helpful discussions.

References

1. Arora, S., "Polynomial-time approximation schemes for Euclidean TSP and other geometric problems," *Proc. 37th IEEE Symp. on Foundations of Computer Science*, IEEE Computer Society, Los Alamitos, CA, 1996, 2–12.
2. Arora, S., "Nearly linear time approximation schemes for Euclidean TSP and other geometric problems," *Proc. 38th IEEE Symp. on Foundations of Computer Science*, IEEE Computer Society, Los Alamitos, CA, 1997, 554–563.
3. Barvinok, A.I., "Two algorithmic results for the traveling salesman problem," *Math. of Oper. Res.* **21** (1996), 65–84.
4. Gusfield, D., Martel, C., and Fernandez-Baca, D., "Fast algorithms for bipartite network flow," *SIAM J. Comput.* **16** (1987), 237-251.
5. Itai, A., Papadimitriou, C., and Swarcfiter, J.L., "Hamilton paths in grid graphs," *SIAM J. Comput.* **11** (1982), 676–686.
6. Lawler, E.L., Lenstra, J.K., Rinnooy Kan, A.H.G., and Shmoys, D.B., *The Traveling Salesman Problem*, Wiley, Chichester, 1985.
7. Megiddo, N., and Tamir, A., "Linear time algorithms for some separable quadratic programming problems," *Oper. Res. Lett.* **13** (1993), 203–211.
8. Mitchell, J., "Guillotine subdivisions approximate polygonal subdivisions: Part II – A simple PTAS for geometric k-MST, TSP, and related problems," *preliminary manuscript*, April 30, 1996.
9. Papadimitriou, C.H., and Yannakakis, M., "The traveling salesman problem with distances one and two," *Math. of Oper. Res.* **18** (1993), 1–11.
10. Trevisan, L., "When Hamming meets Euclid: The approximability of geometric TSP and MST," *Proc. 29th Ann. ACM Symp. on Theory of Computing*, ACM, New York, 1997, 21–29.
11. Zemel, E., "An $O(n)$ algorithm for the linear multiple choice knapsack problem and related problems," *Inf. Proc. Lett.* **18** (1984), 123–128.